arXiv:1609.03340v1 [math.PR] 12 Sep 2016

SHADOW COUPLINGS

MATHIAS BEIGLBÖCK AND NICOLAS JUILLET

Abstract. A classical result of Strassen asserts that given probabilities μ, ν on the real line which are in convex order, there exists a *martingale coupling* with these marginals, i.e. a random vector (X_1, X_2) such that $X_1 \sim \mu, X_2 \sim \nu$ and $\mathbb{E}[X_2|X_1] = X_1$. Remarkably, it is a non trivial problem to construct particular solutions to this problem. In this article, we introduce a family of such martingale couplings, each of which admits several characterizations in terms of optimality properties / geometry of the support set / representation through a Skorokhod embedding. As a particular element of this family we recover the (left-) curtain martingale transport, which has recently been studied [6, 14, 10, 4] and which can be viewed as a martingale analogue of the classical monotone rearrangement. As another canonical element of this family we identify a martingale coupling that resembles the usual *product coupling* and appears as an optimizer in the general transport problem recently introduced by Gozlan et al. In addition, this coupling provides an explicit example of a Lipschitz-kernel, shedding new light on Kellerer's proof of the existence of Markov martingales with specified marginals.

Keywords: Strassen's theorem, Kellerer's theorem, peacocks, (martingale) optimal transport, general transport costs, Skorokhod embedding

1. Introduction

Given Polish spaces X, Y, a measure π on $X \times Y$ with marginals μ and ν is called a *transport plan* from μ to ν or a *coupling* of μ and ν. Let $\Pi(\mu, \nu)$ be the space of transport plans of marginals μ and ν. We will usually consider probability measures μ, ν on the real line having first moments in which case the set of *martingale transport plans* is defined as

$$\Pi_M(\mu, \nu) = \{\pi = \mathrm{Law}(X, Y) \in \Pi(\mu, \nu),\ \mathbb{E}(Y|X) = X\} \tag{1}$$

$$= \{\pi \in \Pi(\mu, \nu) : \int y \, \mathrm{d}\pi_x = y \text{ for } \hat{\mu}\text{-a.e. } (u, x)\ \}. \tag{2}$$

Here the constraint $\mathbb{E}(Y|X) = X$ means that $\mathbb{E}(Y|X = x) = x$ for μ-almost every $x \in \mathbb{R}$, while $(\pi_x)_x$ denotes the disintegration of π wrt μ. By a classical result of Strassen [28], $\Pi_M(\mu, \nu)$ is non-empty if and only if μ, ν are in the convex order $\mu \preceq_C \nu$, i.e. both measures have finite first moments and $\int \phi \, \mathrm{d}\mu \leq \int \phi \, \mathrm{d}\nu$ for every convex $\phi : \mathbb{R} \to \mathbb{R}$.

In [6] we introduced the (left-) curtain coupling π_{lc} which can be seen as a martingale analogue of the monotone rearrangement coupling. An explicit description is provided when μ is finitely supported ([6, Chapter 2]). Another construction using differential equations is given by Henry-Labordère and Touzi [14] for sufficiently regular distributions. According to [19] the coupling method is continuous so that all left-curtain couplings for general measures μ and ν can be approximated using either of the two mentioned constructions, see [19, Remark 2.18].

2010 *Mathematics Subject Classification.* 60G42.

Key words and phrases. couplings, martingales, peacocks, convex order, optimal transport.

The first author acknowledges support through FWF grant Y782. The second author was partially supported by the "Programme ANR JCJC GMT" (ANR 2011 JS01 011 01).

In this paper we will view π_{lc} as one extreme of an infinite family of martingale couplings whose construction is based each on a different parametrization of μ: the curtain coupling will be recovered as one 'end' of this family for a *horizontal* parametrization of μ (a curtain closed from left to right), while at the other end of the spectrum using a *vertical* parametrization we will obtain a new and rather different type of canonical coupling that we shall call sunset coupling π_{sun}. This coupling can be seen as the martingale analogue of the product coupling $\mu \times \nu$ (in classical optimal transport). In view of this it is natural that π_{sun} does not appear as an optimizer of the martingale version of the transport problem. However, we shall see in Theorem 1.1 and Section 5 below that enjoys some optimality properties of a different type. A further particular property is that π_{sun} yields an explicit example of a martingale transport plan which has the *Lipschitz(-Markov) property*. The existence of a martingale transport plan with this property is a key ingredient in all (to the best of our knowledge) proofs of Kellerer's Theorem [21] on the existence of Markov martingales with given marginals. Previous constructions of such transport plans were either non-constructive or relied on particular solutions to the Skorokhod embedding problem.

1.1. **Notation and Main Results.** We write λ for the Lebesgue measure on the unit interval and assume that $\mu \preceq_C \nu$. We fix a *lift* of μ, that is, a probability $\hat{\mu} \in \Pi(\lambda, \mu)$ that will serve as a parameter in the construction of a general version of the left-curtain coupling. The set of *lifted martingale transport plans* is

$$\hat{\Pi}_M(\hat{\mu}, \nu) := \left\{ \hat{\pi} \in \Pi(\hat{\mu}, \nu) : \int y \, \mathrm{d}\hat{\pi}_{u,x} = x \text{ for } \hat{\mu}\text{-a.e. } (u, x) \right\},$$

where $(\pi_{u,x})$ denotes the disintegration of $\hat{\pi}$ wrt $\hat{\mu}$.

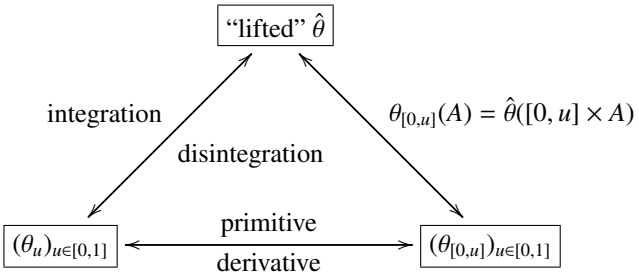

We shall use two further ways to denote the objects $\hat{\mu} \in \Pi(\lambda, \mu) \subseteq \mathcal{P}([0,1] \times \mathbb{R})$ and $\hat{\pi} \in \Pi(\hat{\mu}, \nu) \subseteq \mathcal{P}([0,1] \times \mathbb{R}^2)$: Given a measure $\hat{\theta}$ on $[0,1] \times \mathbb{R}^d$ (where $d = 1, 2$), with $\mathrm{proj}_{[0,1]}(\hat{\theta}) = \lambda$, we write

(1) $(\theta_u)_{u\in[0,1]}$ for the (λ-a.s. unique) disintegration of $\hat{\theta}$ wrt λ.
(2) $(\theta_{[0,u]})_{u\in[0,1]}$ for the family of measures defined for every $u \in [0,1]$ by

$$\theta_{[0,u]}(A) = \hat{\theta}([0,u] \times A) = \int_0^u \theta_s(A)\, ds.$$

Our main result is the following.

Theorem 1.1. *Let μ, ν be real probability measures in convex order and $\hat{\mu} \in \Pi(\lambda, \mu)$. There exists a unique $\hat{\pi} \in \hat{\Pi}_M(\hat{\mu}, \nu)$ satisfying any, and then all of the following properties:*

(1) $\hat{\pi}$ minimizes

$$\hat{\gamma} \mapsto \int (1-u)\sqrt{1+y^2}\, \mathrm{d}\hat{\gamma} \tag{3}$$

on the set $\hat{\Pi}_M(\hat{\mu}, \nu)$.

(2) $\hat{\pi} = \mathrm{Law}(U, B_0, B_\tau)$*, where* (B_t) *is one dimensional Brownian motion,* $\mathrm{Law}(U, B_0) = \hat{\mu}$ *and* τ *is the hitting time of the process* $t \mapsto (U, B_t)$ *into a* left barrier *(i.e. a Borel set* $R \subseteq [0,1] \times \mathbb{R}$ *such that* $(u, x) \in R$, $v \leq u$ *implies* $(v, x) \in R$*).*

(3) $\hat{\pi}(\hat{\Gamma}) = 1$ *for a Borel set* $\hat{\Gamma} \subseteq [0,1] \times \mathbb{R} \times \mathbb{R}$ *which is* monotone *in the sense that for all* $s, t, x, x', y^-, y^+, y'$

$$s < t, (s, x, y^-), (s, x, y^+), (t, x', y') \in \Gamma \Rightarrow y' \notin]y^-, y^+[.$$

(4) *For all* $u \in [0,1]$*, the projection of* $\pi_{[0,u]}$ *onto the second coordinate is the* shadow *of* $\mu_{[0,u]}$ *onto the measure* ν*.*

We add some comments to this result:

- To make sense of the last point, note that if $\mu' \leq \mu$ and $\mu \preceq_C \nu$, then the set $\{\nu' : \mu' \preceq_C \nu' \text{ and } \nu' \leq \nu\}$ is non-empty and has a smallest element $S^\nu(\mu')$ wrt $\preceq_C$, the *shadow* of μ' onto the measure ν (cf. [6, Lemma 4.6] / Definition 2.4 below). Intuitively speaking, among all measures $\nu' \leq \nu$ which are larger than μ' in convex order, $S^\nu(\mu')$ is the most concentrated one.
- We will see below in Proposition 4.3 that $\hat{\pi} \in \hat{\Pi}_M(\hat{\mu}, \nu)$ implies in the setting of (2) that the martingale $(B^\tau_t)_{t \geq 0}$ is uniformly integrable.
- In (1) the cost $c : (u, x, y) \mapsto (1-u)\sqrt{1+y^2}$ can be replaced by any positive $c(u, x, y) = \varphi(u)\psi(y)$ where $\varphi \geq 0$ is strictly decreasing and $\psi \geq 0$ is strictly convex and the minimum over $\hat{\Pi}_M(\hat{\mu}, \nu)$ is finite. More generally the same conclusions holds for a non-negative c with $c_{uyy} < 0$ in a weak sense. Alternative assumptions to $c \geq 0$ are that $\int |\varphi| \, \mathrm{d}\lambda$, $\int |\psi(y)| \, \mathrm{d}\nu < \infty$ or that $c(x, y) \geq A + Bx + Cy$.

We call the unique element of $\hat{\Pi}_M(\hat{\mu}, \nu)$ characterized in Theorem 1.1 the *lifted shadow coupling* (with lift $\hat{\mu}$). Its projection onto the two last coordinates is an element π of $\Pi_M(\mu, \nu)$ that we call *shadow coupling* of μ and ν associated to the lift $\hat{\mu}$. Note that $\pi = \pi_{[0,1]}$ in the terminology introduced in Section 1.1.

If the lift $\hat{\mu}$ is concentrated on the graph of a 1-1 function $T : [0,1] \to \mathbb{R}$ there is an obvious correspondence between elements of $\Pi_M(\mu, \nu)$ and $\hat{\Pi}_M(\hat{\mu}, \nu)$. In particular the optimality property stated in Theorem 1.1 (1) then translates to an optimality property for the martingale version of the transport problem; early papers to investigate such problems include [18, 3, 12, 11, 9, 17, 10, 1, 7]. For general lifts, the shadow coupling does not exhibit particular optimality properties for the martingale transport problem. However, it is characterized by a general optimality problem in the sense of Gozlan et al. [13]. We shall discuss this in Section 5 below.

Canonical choices of lifts lead to canonical martingale couplings of shadow type. We shall be particularly interested in the cases where $\hat{\mu}$ is either the quantile or the product coupling of λ and μ:

- The quantile coupling λ and μ is the unique coupling $\hat{\mu}$ whose support is an increasing function (which then is of course the quantile function of μ). Considering the corresponding lifted shadow coupling in $\hat{\Pi}_M(\hat{\mu}, \nu)$, we recover the *left-curtain coupling* introduced in [6]. We shall henceforth denote this coupling by π_{lc}.

  Notably most of the results established for π_{lc} in [6] are a particular consequence of Theorem 1.1 (cf. Remark 4.6).
- The other shadow coupling we will be particularly interested in is our new *sunset coupling* π_{sun}. It is based on the product lift $\hat{\mu} = \lambda \otimes \mu$, i.e. the independent coupling of λ and μ.

Note that the sources of the above martingale couplings are the two most natural coupling methods for elements in the space without constraint $\Pi(\mu,\nu)$. Looking at the measure μ as the hypograph of its unit density function, we note that the curves $(\mu_{[0,u]})_{u\in[0,1]}$ and $(\mu_u)_{u\in[0,1]}$ reminds a curtain closed from the left in the first case and from the bottom in the second case. This motivates the names *curtain coupling* and *sunset coupling*. The lifted versions are naturally denoted by $\hat\pi_{\mathrm{lc}}$, $\hat\pi_{\mathrm{sun}}$ and called *lifted curtain coupling*, *lifted sunset coupling* respectively.

1.2. **Kellerer's theorem and the sunset coupling.** Concerning the continuous setting, Kellerer's Theorem [22] states that if a family of measures $(\mu_t)_{t\in\mathbb{R}_+}$ satisfies

$$s \le t \Rightarrow \mu_s \preceq_C \mu_t,$$

there exists a martingale $(X_t)_{t\in\mathbb{R}_+}$ with $\mathrm{Law}(X_t) = \mu_t$ for every t. The martingale can be supposed *Markovian* and this is, as far as we are concerned, the most spectacular achievement of this theorem. In contemporary terms (see [15]), $(\mu_t)_{t\in\mathbb{R}_+}$ is called a *peacock* and $(X_t)_{t\in\mathbb{R}_+}$ is a Markovian martingale *associated* to this peacock.

To our best knowledge, all proofs of Kellerer's theorem are based on approximation arguments using sequences of Markov processes. Here the obstacle is that the Markov-property is *not* preserved when passing to the limit. The key insight of Kellerer was to consider Lipschitz-Markov processes, that is Markov-processes whose transition kernels have the following *Lipschitz*-property: a kernel $P : x \mapsto \pi_x$ is called Lipschitz (or more precisely 1-Lipschitz) if $W(\pi_x,\pi_{x'}) \le |x - x'|$ for all x, x' (cf. (5) for the precise definition of the Wasserstein-1 distance). It is then not difficult to see that the property of being a Lipschitz-Markov process *is* preserved when passing to the limit in the sense of finite-dimensional distributions.

Martingale transport plans can be seen as a one step martingales and it is possible to compose several of them for defining a discrete Markovian martingale. The main technical step in Kellerer's proof is therefore to show that given measures μ, ν in convex order there exists a martingale transport plan which has the Lipschitz property.

Let us add that Lipschitz kernels for measures in convex orders do not exist in higher dimensions $d \ge 2$ (see for instance [20]). As Lipschitz kernels and their avatars are the only known methods for proving the Kellerer theorem, still to our best knowledge, it is an open problem whether this theorem holds in dimensions greater than or equal to two.

While non of the various extremal martingale couplings constructed in [18, 6, 17, 10, 27] has the Lipschitz property, we shall see that the sunset coupling has the Lipschitz-property. Moreover, it connects to Kellerer's original proof (in [21]) of the existence of Lipschitz kernels which we now describe:

In Kellerer's terminology ([21, 22]), martingale transport plans π appear as pairs consisting of an initial measure μ together with a transition kernel $P : x \mapsto \pi_x$, satisfying $\int y\,\mathrm{d}\pi_x(y) = x$ and $\mu P = \nu$. Slightly abusing notation, we will occasionally identify martingale transport plans with their kernels. The strategy of Kellerer in order to find a Lipschitz kernel is to use Choquet's theory for describing ν as a combination of extreme measures of $E(\mu)$, the set of measures greater than μ in the convex order. He proves that if ω is an extreme element, the set $\Pi_M(\mu,\omega)$ consists only of a single element π and this element π has the Lipschitz property (see Proposition 2.2). Using that the Lipschitz-property is preserved under convex combinations, Kellerer then obtains that $\Pi_M(\mu,\nu)$ contains some Lipschitz kernel.

Let us now discuss how $\pi_{\mathrm{sun}} \in \Pi_M(\mu,\nu)$ compares to Kellerer's proof of the existence of Lipschitz kernels. In fact according to the barrier characterization in Theorem 1.1, we

have

$$\pi_{\text{sun}} = \int_0^1 \hat{\pi}_u \, \mathrm{d}u.$$

Here $\hat{\pi}_u = \text{Law}(B_0, B_\tau \mid U = u)$ where (B_t) is a Brownian motion with starting distribution μ, and τ (conditioned on $U = u$) is the hitting time of the barrier's vertical section $R_u := \{y \in \mathbb{R} : (u, y) \in R\}$, see Figure 1. We denote by μP_{R_u} the hitting law $\text{Law}(B_\tau \mid U = u)$ and call the corresponding kernel P_{R_u} the *hitting projection* onto R_u (see also Definition 2.1). We have $\nu = \int_0^1 \mu P_{R_u} \, \mathrm{d}u$.

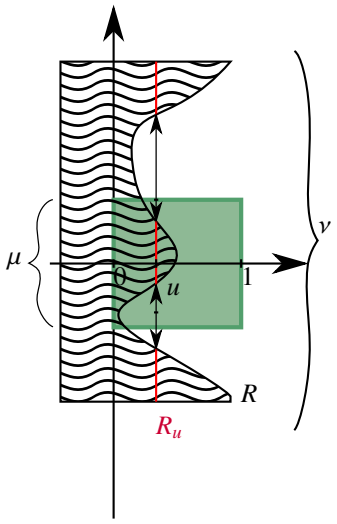

FIGURE 1. Hitting the barrier

Kellerer established that the measures $\omega = \mu P_T$ obtained through a hitting projection are exactly the extreme elements of $E(\mu)$ (provided one obtains a *martingale* transport plan, that is $\text{spt}(\mu) \subseteq [\inf(T), \sup(T)]$). Recalling the argument two paragraphs above, $\hat{\pi}_u = \mu(\text{Id} \times P_{R_u})$ is the unique element of $\Pi_M(\mu, \mu P_{R_u})$, the kernels P_{R_u} are Lipschitz kernels, and in particular, the sunset coupling has the Lipschitz property.

We stress that the uniqueness in Theorem 1.1 permits us to gurantee that there exists a *unique* Choquet representation through a family $(R_u)_{u \in [0,1]}$ which is *ordered* in the sense that $s \leq t$ implies that $R_t \supseteq R_u$. Indeed, the definition of barrier tells us that $R_v \subseteq R_u$ if $u \leq v$ so that such a representation exists. If some Choquet representation of ν is given by measures ν_u obtained using the hitting projection on sets R_u which is ordered in the above sense, then these sets consitute a barrier and based on the uniqueness assertion in Theorem 1.1, the family $(\nu_u)_{u \in [0,1]}$ with $\nu = \mu P_{R_u}$ is the one associated to the sunset coupling.

Let us summarize this in a theorem.

Theorem 1.2 (Sunset coupling and Lipschitz kernel)**.** *Let μ and $\nu \in \mathcal{P}$ satisfy $\mu \preceq_C \nu$. Then there exists a probability measure χ on $\mathcal{I} = \{T \subseteq \mathbb{R} : T \text{ is closed}, \inf T = -\infty, \sup T = +\infty\}$ which represents ν in the sense that*

$$\nu(A) = \textstyle\int_{\mathcal{I}} \mu P_T(A) \, \mathrm{d}\chi(T) =: \mu P_\chi(A) = \mu_\chi(A). \tag{4}$$

Moreover P_χ is a Lipschitz kernel and $\mu(\text{Id} \times P_\chi) \in \Pi_M(\mu, \nu)$ is a Lipschitz martingale plan.

A possible choice of χ is the uniform measure on $(R_u)_{u\in[0,1]}$ where R_u is a decreasing family for $\subseteq$. Moreover if χ' is another measure associated to a non-increasing family $(R'_u)_{u\in[0,1]}$ with $\nu = \mu P_{\chi'}$ and $\mu(\mathrm{Id}\times P_{\chi'}) \in \Pi_M(\mu,\nu)$, then $\mu(\mathrm{Id}\times P_\chi) = \mu(\mathrm{Id}\times P_{\chi'})$. (This measure is the sunset coupling of μ and ν, that is the shadow coupling of μ with lift $\lambda\times\mu$ and ν.)

Another proof that $\Pi_M(\mu,\nu)$ contains at least one Lipschitz kernel is given by solutions of the Skorohod embedding problem. Root's embedding is considered in [5] and Hobson's embedding in [23, Lemma 3.3]. To our best knowledge, this is after Kellerer's original proof the second known type of proof. Theorem 1.1 finally spans a bridge between the two methods because it presents a kind of lifted Skorohod embedding.

2. Preliminaries

2.1. **Definitions, notations on the martingale transport problem.** We consider the space $\mathcal{M}$ of positive measures on $\mathbb{R}$ with finite first moments. The subspace of probability measures with finite expectations is denoted by $\mathcal{P}$. In higher dimensions we denote the corresponding spaces by $\mathcal{M}(\mathbb{R}^d)$ and $\mathcal{P}(\mathbb{R}^d)$. For $\mu,\ \nu\in\mathcal{M}$, the Wasserstein-1 distance is defined by

$$W(\mu,\nu) = \sup_{f\in\mathrm{Lip}(1)} \left| \int f\,\mathrm{d}\mu - \int f\,\mathrm{d}\nu \right| \tag{5}$$

and endows $(\mathcal{P}, W)$ with $\mathcal{T}_1$, the usual topology for probability measures with finite first moments. In the definition, the supremum is taken over all 1-Lipschitzian functions $f : \mathbb{R}\to\mathbb{R}$. We also consider W (with the same definition) on the subspace $m\mathcal{P} = \{\mu\in\mathcal{M} | \mu(\mathbb{R}) = m\}\subseteq\mathcal{M}$ of measures of mass m.

According to the Kantorovich duality theorem, an alternative definition in the case $\mu,\nu\in\mathcal{P}$ is

$$\inf_{(\Omega,X,Y)} \mathbb{E}(|Y-X|) \tag{6}$$

where $X,\ Y : (\Omega,\mathcal{F},\mathbb{P})\to\mathbb{R}$ are random variables with marginals μ and ν. The infimum is taken among all joint laws (X,Y), the probability space $(\Omega,\mathcal{F},\mathbb{P})$ being part of the minimisation problem. Note that without loss of generality $(\Omega,\mathcal{F},\mathbb{P})$ can be assumed to be $([0,1],\mathcal{B},\lambda)$ where λ is the Lebesgue measure and $\mathcal{B}$ the σ-algebra of Borel sets on $[0,1]$.

A special choice of 1-Lipschitzian function is the function $f_t : x\in\mathbb{R}\to|x-t|\in\mathbb{R}$. Therefore if $\mu_n\to\mu$ in $\mathcal{M}$, the sequence of functions $u_{\mu_n} : t\mapsto\int f_t(x)\,\mathrm{d}\mu_n(x)$ converges to u_μ pointwise. The converse statement also holds if all the measures have the same mass and barycenter (see [16, Proposition 2.3] or [6, Proposition 4.2]). For $\mu\in\mathcal{M}$, the function u_μ is usually called the *potential function* of μ.

2.2. **Bijection between curves, primitive curves, and lifted measures.** We elaborate on the equivalent avatars of the lifted measures introduced in Subsection 1.1. In short, we are representing the same mathematical object in three ways: we consider the measure $\hat\theta$, the almost surely defined disintegration $(\theta_u)_{u\in[0,1]}$, and the primitive curve $(\theta_{[0,u]})_{u\in[0,1]}$. We first recall the integrability conditions. The lifted measure $\hat\theta$ is a probability measure on $[0,1]\times\mathbb{R}^d$ (where $d=1$ or $d=2$) such that $\hat\theta(\hat\rho) < +\infty$ is finite where $\rho : x\mapsto\|x\|_{\mathbb{R}^d}$ and $\hat\rho(u,x) = \rho(x)$. This integrability condition corresponds to $\theta_{[0,1]}(\rho) < +\infty$ for the primitive curve $(\theta_{[0,u]})_{u\in[0,1]}$ and $\int_0^1\theta_u(\rho)\,\mathrm{d}u < +\infty$ for $(\theta_u)_{u\in[0,1]}$. The marginal condition asserts that $\hat\theta\in\Pi(\lambda,\theta)$ for some $\theta\in\mathcal{P}(\mathbb{R}^d)$. In terms of the primitive curve this can be expressed by asserting that $\theta_{[0,1]} = \theta$ and $\theta_{[0,u]}(\mathbb{R}) = u$. The equivalent condition on $(\theta_u)_{u\in[0,1]}$ is

that λ-almost surely $\theta_u \in \mathcal{P}$ and $\theta = \int_0^1 \theta_u \, \mathrm{d}u$. Note that from a probabilistic point of view, if (U, X) is a random vector of law $\hat{\theta}$ with $U \sim \lambda$, the other representations are given by $(u \times \mathrm{Law}(X|U \leq u))_{u\in[0,1]}$ and $(\mathrm{Law}(X|\, U = u))_{u\in[0,1]}$. Finally the object that we will Ultimately be most interested in is not the lift $\hat{\theta}$ but $\theta := \theta_{[0,1]} = \mathrm{Law}(X)$.

In what follows we explain that the derivative of the primitive curve $(\theta_{[0,u]})_{u\in[0,1]}$ can be considered with respect to $\mathcal{T}_1$, the weak topology wrt to continuous functions which have at most linear growth. Let us start with $\hat{\theta} \in \Pi(\lambda, \theta)$. We disintegrate the measure with respect to the first marginal and obtain an a.s. uniquely determined family $(\theta_u)_{u\in[0,1]}$ such that for almost every u, $\theta_u \in \mathcal{P}(\mathbb{R}^d)$. (We can assume that the measure is zero for the other parameters.) Define $\theta_{[0,u]}$ for $u \in [0, 1]$ by

$$\theta_{[0,u]}(A) = \hat{\theta}([0, u] \times A) = \int_0^u \theta_s(A) \, \mathrm{d}s$$

for $A \subseteq \mathbb{R}^d$ Borel. Given a function $f : x \in \mathbb{R}^d \to \mathbb{R}$ with $f(x)/(1 + \|x\|)$ bounded, the function $s \mapsto \theta_s(f)$ is measurable and in $L^1([0, 1])$. Hence at almost every time $u \in [0, 1]$ the function $t \mapsto \theta_{[0,u]}(f) = \int_0^u \theta_t(f) \, \mathrm{d}t$ is differentiable with derivative $\theta_u(f)$. It is important that the set $L \subseteq [0, 1]$ of times at which the derivative exists for all f is a Borel set of full measure, as we will verify in the next paragraph. Before establishing this claim, note that this permits us to define a canonical disintegration $(\tilde{\theta}_u)_{u\in[0,1]}$: we define the measure $\tilde{\theta}_u$ as the derivative if $u \in L$, and zero otherwise.

We turn now to the claim: Let $\mathcal{X}$ be a countable set of functions which is dense in the space $C_c(\mathbb{R}^d)$ of functions with compact support and let $\mathcal{X}_+$ be $\mathcal{X} \cup \{\rho\}$ where $\rho(x) = \|x\|_{\mathbb{R}^d}$. Let $L \subseteq [0, 1]$ be the set such that at any time $u \in L$, θ_u is a probability measure and $u \mapsto \theta_{[0,u]}(f)$ has derivative $\theta_u(f)$ for any $f \in \mathcal{X}_+$ and note that L has full mass. Then, as an increment h goes to zero the measure $h^{-1}(\theta_{[0,u+h]} - \theta_{[0,u]})$ weakly converge to θ_u and as $\rho \in \mathcal{T}_1$ convergences holds also in $\mathcal{T}_1$, cf. [29, Theorem 7.12]. Thus L is a set of differentiation for any function with finite first moment.

2.3. **General description of the construction.** In what follows we shortly explain the canonical scheme to define $\hat{\pi} \in \hat{\Pi}_M(\hat{\mu}, \nu)$ where the marginals $\hat{\mu} \in \Pi(\lambda, \mu)$ and $\nu \in \mathcal{P}$ are given. Recall that the resulting coupling $\pi = \pi_{[0,1]} = (\mathrm{proj}_{x,y})_{\#}\hat{\pi}$ of $\mu = \mu_{[0,1]}$ and ν fits more naturally to the theory of optimal transportation than the lifted coupling $\hat{\pi}$.

Represent $\hat{\mu}$ in the form $(\mu_{[0,u]})_{u\in[0,1]}$. The first canonical operation, called *shadow projection* on ν, consists in building the curve $(\nu_{[0,u]})_{u\in[0,1]}$ from it (see Defintion 2.4). Hence the construction is complete if on a set $L \subseteq [0, 1]$ of differentiation (of full measure) of $(\mu_{[0,u]})_u$ and $(\nu_{[0,u]})_u$ we know how to canonically choose a joint law π_u of μ_u and ν_u. In our situation, the martingale constraint and the fact that we used the shadow projection will make this choice uniquely determined and related to Kellerer's *hitting projection* (see Definition 2.1). We will thus obtain $(\pi_u)_{u\in[0,1]}$ and equivalently $(\pi_{[0,u]})_{u\in[0,1]}$ and $\hat{\pi}$. This construction will be carried out in detail in the proof of Theorem 3.1.

2.4. **More details on Kellerer's approach.** In relation with Theorem 1.2 let us present in a more formal way the Choquet representation given in [21, 22].

We denote by $\mathcal{F}(\mathbb{R})$ the space of closed subsets of $\mathbb{R}$, and $\mathcal{I}$ the subspace of those elements $T \in \mathcal{F}(\mathbb{R})$ such that $\sup T = -\inf T = +\infty$. The space $\mathcal{F}(\mathbb{R})$ is endowed with the coarsest topology for which the mappings $T \in \mathcal{F}(\mathbb{R}) \mapsto d(x, T) \in [0, \infty]$ are continuous. By [22, Satz 13] this topology is metrisable and compact, and by [22, Satz 14] $\mathcal{I}$ is a G_δ subset of $\mathcal{F}(\mathbb{R})$.

Definition 2.1 (Hitting projection of measure in/to a set)**.** Let T be an element of $\mathcal{I}$. For every $x \in T$, let $x_T^- = \sup(T \cap (-\infty, x])$ and $x_T^+ = \inf(T \cap [x, +\infty))$. We define now the

Kellerer dilation ([22, Definition 16])

$$P_T(x,\cdot) = \begin{cases} \delta_x & \text{if } x \in T; \\ (x_T^+ - x_T^-)^{-1}[(x_T^+ - x)\delta_{x_T^-} + (x - x_T^-)\delta_{x_T^+}] & \text{otherwise.} \end{cases}$$

Hence if $\mu \in \mathcal{P}$, the hitting projection of μ in T is $\nu = \mu P_T$ and the hitting coupling of μ and ν is given by $\pi(A \times B) = \int_A P_T(x, B)\,\mathrm{d}\mu(x)$; we shall abbreviate this by $\pi = \mu(\mathrm{Id} \times P_T)$.

Note that if T is not an element of $\mathcal{I}$ but $\mathrm{spt}(\mu) \subseteq [\inf T, \sup T]$, the hitting projection of μ still makes sense because x_T^- and x_T^+ are finite and thus π is a martingale transport plan. It is easy to replace T by an element $T^* \in \mathcal{I}$ such that the resulting hitting couplings are equal, simply define $T^* =]-\infty, \inf T] \cup \bar{T} \cup [\sup(T), +\infty[$. Note that if $\mu \preceq_C \nu$, this remark applies to the set $T = \mathrm{spt}(\nu)$ since $\mathrm{spt}(\mu) \subseteq [\inf T, \sup T]$ is satisfied.

Proposition 2.2. *Let μ be an element of $\mathcal{P}$ and $T \in \mathcal{F}(\mathbb{R})$ satisfy $\mathrm{spt}(\mu) \subseteq [\inf T, \sup T]$. Let ν be μP_T. Then, the hitting coupling $\mu(\mathrm{Id} \times P_T)$ is the unique element of $\Pi_M(\mu, \nu)$.*

Proof. This is [22, Satz 25]. Alternatively, one may consider the decomposition of (μ, ν) into irreducible components described in [6, Theorem 8.4]. This theorem specifies two canonical series $\mu = \mu \wedge \nu + \sum \mu_n$ and $\nu = \mu \wedge \nu + \sum \nu_n$ such that any $\pi \in \Pi_M(\mu, \nu)$ can be written as $\pi = (\mathrm{Id} \times \mathrm{Id})_\#(\mu \wedge \nu) + \sum \pi_n$ with $\pi_n \in \Pi_M(\mu_n, \nu_n)$ for every n. This decomposition is based on the potential functions u_μ and u_ν: the set $\{u_\mu < u_\nu\}$ is open and hence consists of a (finite or countable) union of open intervals I_n. Then $\mu_n = \mu|_{I_n}$ whereas ν is concentrated on $\bar{I}_n$. In the present situation, u_ν is affine on each interval $\bar{I}_n$. This implies that necessarily $\pi_n = \mu_n(\mathrm{Id} \times P_{\bar{I}_n}) = \mu_n(\mathrm{Id} \times P_T)$ and hence $\pi = \mu(\mathrm{Id} \times P_T)$. □

Let $\mathcal{U}$ be the space of probability measures χ on $\mathcal{I}$ such that $\int_{\mathcal{I}} d(0,T)\,\mathrm{d}\chi(T)$ is finite, endowed with the coarsest topology such that $\chi \mapsto \int_{\mathcal{I}} h(T)\,\mathrm{d}\chi(T)$ is continuous for every continuous functions $h : \mathcal{I} \mapsto \mathbb{R}$ with $\sup_T h(T)(1 + d(0,T))^{-1} < +\infty$. [22, Satz 18] asserts that $\mathcal{U}$ is metrisable and [22, Satz 19] that $\mathcal{U}_0 = \{\chi \in \mathcal{U} | \chi - \text{almost surely } T \supseteq T_0\}$ is compact. Moreover, according to [22, Satz 20], for $\mu \in \mathcal{P}$, the measure $\nu = \mu P_\chi$ has also finite first moments.

Kellerer establishes in [22, Theorem 1] that for a given $\mu \in \mathcal{P}$ the extreme points of $\{\nu \in \mathcal{P} : \mu \preceq_C \nu\}$ exactly are the measures μP_T. The latter set is not compact but going first through the compact and convex sets $\{\nu \in \mathcal{P},\ \mu \preceq_C \nu \preceq_C \mu_S\}$ for sets $S \in \mathcal{I}$, Kellerer is able to derive a Choquet representation in [22, Theorem 4].

Theorem 2.3 (A Choquet representation established by Kellerer)**.** *Let μ and $\nu \in \mathcal{P}$ satisfy $\mu \preceq_C \nu$. Then there exist a probability measure $\chi \in \mathcal{U}$ such that*

$$\nu = \int_{\mathcal{I}} \mu P_T\,\mathrm{d}\chi(T) =: \mu P_\chi = \mu_\chi.$$

Let us give some precisions on this result and on Theorem 1.2.

- Theorem 1.2 improves Theorem 2.3 with a uniqueness statement based on the natural order on $\mathcal{I}$. Our proof is independent of Kellerer's. In particular we do not study what the extreme elements of $E(\mu) = \{\omega \in \mathcal{P} | \mu \preceq_C \omega\}$ are.
- As explained in Subsection 1.2 and using the notation of Subsection 2.3, for every $u \in [0,1]$ the transport plan π_u transfers $\mu_u = \mu$ onto ν_u. We will see in Proposition 2.7 that this transport is given through the *hitting projection* of μ onto $\mathrm{spt}(\nu - \nu_{[0,u]})$. This description also provides a canonical choice for R_u in Theorem 1.2 as well as a canonical barrier in Theorem 1.1 (2) by setting $R = (]-\infty, 0[\times \mathbb{R}) \cup \{(u,y) \in [0,1] \times \mathbb{R},\ y \in \mathrm{spt}(\nu - \nu_{[0,u]})\}$.

- In [22], uniqueness of a measure on the extreme elements of $\{\omega \in \mathcal{P}, \mu \preceq_C \omega\}$ is not claimed and can easily be disproved. Set for instance $\mu = \delta_0$ and $\nu = \frac{\delta_{-2}+\delta_{-1}+\delta_1+\delta_2}{4}$. Taking the uniform probability measure on $T_1 = \mathbb{R}\backslash]-1,1[$ and $T_2 = \mathbb{R}\backslash]-2,2[$ on the one hand and an atomic measure on $\mathbb{R}\backslash]-1,2[, \mathbb{R}\backslash]-2,1[$ and $\mathbb{R}\backslash]-2,2[$ we obtain two different representations of ν. Another, more trivial, type of non-uniqueness can be observed: In the previous example T_1 can also be replaced for instance by $\mathbb{R} \setminus (]-1,1[\cup]-20,-15[)$ providing the same measure on $\mathcal{P}$ but another measure on $\mathcal{I}$.
- From Theorem 1.2 it is not evident that χ satisfies the integrability condition appearing in the definition of $\mathcal{U}$. However according to Lemma 15 in [22], this holds if and only if $\nu = \mu_\chi$ has finite first moments, i.e. is an element of $\mathcal{P}$. But this is one hypotheses in Theorem 1.2. Hence, this theorem completely generalizes Theorem 2.3, also with respect to $\chi \in \mathcal{U}$.

2.5. **Order relations and shadows.** On $\mathcal{M}$ we write $\mu \preceq_{C,+} \nu$ if and only if there exists $\eta \in \mathcal{M}$ with $\mu \preceq_C \eta$ and $\eta \preceq_+ \nu$. Here $\preceq_+$ means $\eta(A) \leq \nu(A)$ for every Borel set A. The order $\preceq_{C,+}$ can also be characterized by asserting $\mu(f) \leq \nu(f)$ for every convex positive function f. We also introduce the stochastic order $\mu \preceq_{\text{sto}} \nu$ that holds if $\mu(f) \leq \nu(f)$ for every integrable increasing function. This is equivalent to $G_\mu \leq G_\nu$ where G_μ denotes the unique increasing left-continuous function with $(G_\mu)_\#\lambda = \mu$, i.e. the quantile function. See [19] for more details in the context of martingale optimal transport.

Definition 2.4 (Definition of the shadow)**.** If $\mu \preceq_{C,+} \nu$ there exists a unique measure η such that

- $\mu \preceq_C \eta$
- $\eta \preceq_+ \nu$
- If η' satisfies the two first conditions (i.e. $\mu \preceq_C \eta' \preceq_+ \nu$), one has $\eta \preceq_C \eta'$.

This measure η is called the *shadow* of μ in ν and we denote it by $S^\nu(\mu)$.

Shadows are sometimes difficult to determine. An important fact is that they have the smallest variance among the set of measures η'. Indeed, $\eta \preceq_C \eta'$ implies $\int x\,\mathrm{d}\eta = \int x\,\mathrm{d}\eta'$ and $\int x^2\mathrm{d}\eta \leq \int x^2\mathrm{d}\eta'$ with equality if and only if $\eta = \eta'$ or $\int x^2\mathrm{d}\eta = +\infty$.

Example 2.5 (Shadow of an atom, Example 4.7 in [6]). Let δ be an atom of mass α at a point x. Assume that $\delta \preceq_{C,+} \nu$. Then $S^\nu(\delta)$ is the restriction of ν between two quantiles, more precisely it is $\nu' = (G_\nu)_\#\lambda_{]s;s'[}$ where $s' - s = \alpha$ and the barycenter of ν' is x.

The following result is one of the most important on the structure of shadows (Theorem 4.8 of [6]).

Proposition 2.6 (Structure of shadows)**.** *Let γ_1, γ_2 and ν be elements of $\mathcal{M}$ and assume that $\mu = \gamma_1 + \gamma_2 \preceq_{C,+} \nu$. Then we have $\gamma_2 \preceq_{C,+} \nu - S^\nu(\gamma_1)$ and*

$$S^\nu(\gamma_1 + \gamma_2) = S^\nu(\gamma_1) + S^{\nu - S^\nu(\gamma_1)}(\gamma_2).$$

An important consequence is that if $(\mu_{[0,u]})_u$ is a primitive curve and $\mu_{[0,1]} \preceq_C \nu$, then the curve $(\nu_{[0,u]})_u$ satisfies $\nu_{[0,u]}(\mathbb{R}) = u$ and using $\gamma_1 = \mu_{[0,u]}$ and $\gamma_1 + \gamma_2 = \mu_{[0,v]}$ we obtain $\nu_{[0,u]} \preceq_+ \nu_{[0,v]}$ for every $u \leq v$. Hence $(\nu_{[0,u]})_u$ is a primitive curve.

We consider the derivatives of shadow curves associated to a primitive curve.

Proposition 2.7. *For $\hat{\mu} \in \Pi(\lambda, \mu)$ and $\mu \preceq_C \nu$, let $u \mapsto \mu_{[0,u]}$ have right derivative μ_{u_0} at u_0 and let $\nu_{[0,u]}$ be $S^\nu(\mu_{[0,u]})$. Then $(\nu_{[0,u]})$ has a right derivative at u_0. This derivative is given by $\mu_{u_0}P_T$ and $\mathrm{spt}(\mu_{u_0}) \subseteq [\inf T, \sup T]$ where T is the support of $\nu_{]u_0,1]} := \nu - \nu_{[0,u_0]}$.*

Consider $h^{-1}(\nu_{[0,u_0+h]} - \nu_{[0,u_0]}) = h^{-1}(S^{\nu}(\mu_{[0,u_0+h]}) - S^{\nu}(\mu_{[0,u_0]})) =: \sigma_h$. According to Proposition 2.6 σ_h equals

$$h^{-1} S^{\nu_{]u_0,1]}}(\mu_{]u_0,u_0+h]})$$

where we set $\mu_{]u,v]} = \mu_{[0,v]} - \mu_{[0,u]}$. But we know that $h^{-1}(\mu_{[0,u_0+h]} - \mu_{[0,u_0]}) = h^{-1}\mu_{]u_0,u_0+h]}$ tends to μ_{u_0} as $h \downarrow 0$. An easy scaling analysis shows that σ_h can in fact be written

$$\sigma_h = S^{h^{-1}\nu_{]u_0,1]}}(h^{-1}\mu_{]u_0,u_0+h]}).$$

We are roughly speaking considering the shadow projection of μ_{u_0} into the infinite measure $\infty \cdot \nu_{]u_0,1]}$. As $\mathcal{T}_1$ is metric, it is enough for the convergence on the right of u_0 that we prove the following lemma:

Lemma 2.8. *Let $\eta_n \to \eta$ in $\mathcal{P}$ and $H_n \to \infty$ and assume for every $n \geq 1$, $\eta_n \preceq_{C,+} H_n\upsilon$. Then $S^{H_n\upsilon}(\eta_n) \to \eta P_T$ in $\mathcal{P}$ and $\mathrm{spt}(\eta) \subseteq [\inf T, \sup T]$ where $T = \mathrm{spt}(\upsilon)$.*

Proof. Note first that due to the convex order relation $\eta_n \preceq_{C,+} \upsilon$ we have $\mathrm{spt}(\eta_n) \subseteq [\inf T, \sup T]$. Going to the limit $\mathrm{spt}(\eta) \subseteq [\inf T, \sup T]$ as well. We are left with the proof of $\eta_n \preceq_{C,+} H_n\upsilon$.

1. We first prove the result if $\eta_n = \delta_x$ for every $n \in \mathbb{N}$. We prove in fact a little stronger statement: if γ_n has mass less than or equal to one and $\gamma_n \preceq_+ H_n\upsilon$, the sequence $S^{H_n\upsilon - \gamma_n}(\delta_x)$ converges to ηP_T. Moreover $x \in T^\circ$, $x \in T^c$ or $x \in \partial T$. In these three cases the result easily follow from Example 2.5.

2. We assume now that for every $n \in \mathbb{N}$, we have $\eta_n = \eta = \sum_{k=1}^n a_k \delta_{x_k}$. The proof is an induction. The initial value $n = 1$ has been previously done. We assume the statement for $n-1 \geq 1$ and prove it for n by using the decomposition $\eta = \eta' + a_n\delta_n$ where $\eta' = \sum_{k=1}^{n-1} a_k\delta_{x_k}$. We have [6, Theorem 4.8]

$$S^{H_n\upsilon}(\eta' + a_n\delta_{x_n}) = S^{H_n\upsilon}(\eta') + S^{\beta_n}(a_n\delta_{x_n})$$

where $\beta_n = H_n\upsilon - S^{H_n\upsilon}(\eta')$. Each of the two terms converges to the Kellerer projection of η' resp. $a_n\delta_n$ onto T. Note that for the second projection we used the full strength of the statement proved in 1.

3. A general measure η can be approximated using a convex combination of Dirac masses η_k with $\eta_k \preceq_C \eta$ and such that $\eta_k \to \eta$ [19, point 3. in proof of Proposition 2.34]. We have

$$W(S^{H_n\upsilon}(\eta_k), S^{H_n\upsilon}(\eta)) \leq W(\eta_k, \eta).$$

This goes to zero uniformly in n as k goes to infinity. But $S^{H_n\upsilon}(\eta_k) \to \eta_k P_T$ as n tends to infinity and the composition with P_T is continuous (cf. [22, Section 2.2], this can be understood easily from the action of P_T on the potential functions). Hence we obtain the result for any constant sequence $\eta_n = \eta$.

4. If η_n is a non-constant sequence

$$W(S^{H_n\upsilon}(\eta_n), \eta P_T) \leq W(S^{H_n\upsilon}(\eta_n), S^{H_n\upsilon}(\eta)) + W(S^{H_n\upsilon}(\eta), \eta P_T),$$

which tends to zero as required ([19, Proposition 2.34]). □

3. Construction

Based on the preparations in the previous section we can now rigorously introduce the *lefted shadow couplings*.

Theorem 3.1 (Existence, construction, and uniqueness of the lifted shadow coupling)**.** *Let μ and ν be elements of $\mathcal{P}$ with $\mu \preceq_C \nu$. Let $\hat{\mu}$ be an element of $\Pi(\lambda, \mu)$. Then there exists a unique element $\hat{\pi} \in \hat{\Pi}_M(\hat{\mu}, \nu)$, the lifted shadow coupling of $\hat{\mu}$ and ν, such that for every*

$u \in [0,1]$, the marginals of $\pi_{[0,u]}$ are $\mu_{[0,u]}$ and it shadow projection $S^\nu(\mu_{[0,u]})$. Moreover $\nu_u = \mu_u P_{T(u)}$ where $T(u)$ is the support of $\nu_{]u,1]} := \nu_{[0,1]} - \nu_{[0,u]}$.

Note that Theorem 3.1 implies in particular that there exists a unique $\hat\pi$ satisfying Theorem 1.1 (4).

Proof of Theorem 3.1. Let $\nu_{[0,u]} = S^\nu(\mu_{[0,u]})$ be as in the statement and let $(\mu_u)_{u\in[0,1]}$ and $(\nu_u)_{u\in[0,1]}$ be the derivative curves. According to Proposition 2.7, ν_u can for almost every $u \in [0,1]$ be identified with the Kellerer projection of μ_u in $T(u) := \mathrm{spt}(\nu_{]u,1]})$ where $\nu_{]u,1]} = \nu - \nu_{]0,u]}$. Hence we can define $\pi_u = \mu_u(\mathrm{Id}\times P_{T(u)}) \in \Pi_M(\mu_u,\nu_u)$ for almost every u, and then the corresponding $\hat\pi$ and $\pi_{[0,u]} = \int_0^u \pi_t\,\mathrm{d}t$.

Conversely, let $\hat\pi$ and $(\pi_{[0,u]})_{u\in[0,1]}$ be as in the statement. The curve has marginals $\mu_{[0,u]}$ and $\nu_{[0,u]} = S^\nu(\mu_{[0,u]})$ so that one can apply Proposition 2.7. At points u where the derivatives π_u, μ_u and ν_u exist, we have $\pi_u \in \Pi_M(\mu_u,\nu_u)$ and $\nu_{[0,u]} = \mu_u P_{T(u)}$ as in the last paragraph, proving the uniqueness part of the statement. □

3.1. **Examples of shadow couplings.** We now further discuss canonical classes of families $(\mu_{]0,u]})_{u\in]0,1]}$ giving rise to particular shadow couplings. We consider different canonical lifts $\hat\mu \in \Pi(\lambda,\mu)$ of μ. The first one is the monotone coupling of λ and μ and the second the independent coupling $\lambda\times\mu$.

- $\mu_{[0,u]} = (G_\mu)_\#\lambda|_{[0,u]}$ (corresponding to the left-curtain coupling);
- $\mu_{[0,u]} = u\cdot\mu$ (corresponding to the sunset coupling);
- $\mu_{[0,u]} = S^\mu(u.\delta_m)$ where $m = \int x\,\mathrm{d}\mu(x)$ (corresponding to the middle curtain coupling). Recall Example 2.5 for the shadow of an atom.

Figures 2, 3 and 4 illustrate the corresponding Skorokhod embeddings when μ and ν are two uniform laws.

3.1.1. *The left- and right-curtain couplings.* This case corresponds to the construction given in [6], even though the construction described there appears slightly different. In fact for $u = F_\mu(x)$ the three marginals of $\hat\pi$ are $\lambda|_{[0,u]}$, $\mu|_{[0,x]}$ and $S^\nu(\mu_{[0,x]})$ so that for every $x\in\mathbb{R}$, π has marginals $\mu|_{[0,x]}$ and $S^\nu(\mu_{[0,x]})$. In [6] this was used to define the left-curtain coupling π_{lc}.

In an entirely symmetric fashion we can define the right-curtain coupling through $\mu_{[0,u]} = (G_\mu)_\#\lambda|_{[1-u,1]}$ for $u\in[0,1]$.

3.1.2. *The sunset coupling.* In this case we have $\mu_u = \mu$ for almost every $u\in[0,1]$ and $\nu_u = \mu P_{T(u)}$ where $T(t) = \mathrm{spt}(\nu - S^\nu(u\cdot\mu))$. Hence

$$\nu = \int_0^1 \mu P_{T(ut)}\,\mathrm{d}u$$

Of course $T(u)$ can be replaced by $T^*(u) = T(u)\cup(]-\infty,\inf\nu]\cup[\sup\nu,+\infty[) \in \mathcal{I}$. This is a possibility to make not only the almost-everywhere defined family $\nu_u = \mu_u P_{T(u)}$ unique as in Theorem 1.2 but also the measure χ.

3.1.3. *The middle-curtain coupling.* As variant of the (left-) curtain coupling, we introduce a middle-curtain coupling. Under the condition that μ and ν are in diatomic order ($\mu \preceq_{DC} \nu$), i.e.

$$S^\mu(u\delta_m) \preceq_C S^\nu(u\delta_m) \quad \text{for every } u\le 1, \tag{7}$$

where m denotes the center of μ, the middle-curtain coupling coincide with an exceptionaly simple martingale transport plan that has been introduced in [20, Section 4]. (Clearly, the diatomic order is more restrictive than the convex order.)

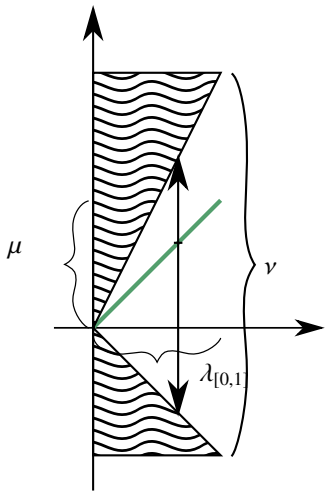

FIGURE 2. Left-curtain coupling of uniform measures

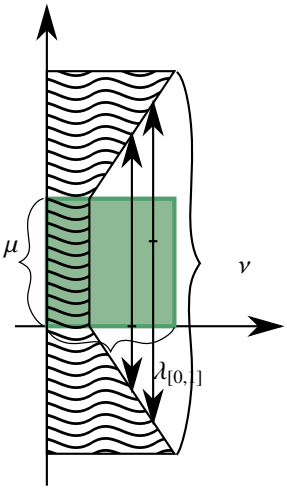

FIGURE 3. Sunset coupling of uniform measures

We define the middle-curtain coupling as the shadow coupling corresponding to the family $\mu_{[0,u]} = S^{\mu}(u\delta_m), u \in [0,1]$. Note that this corresponds to $\mu_u = a(u)\delta_{f(u)} + b(u)\delta_{g(u)}$ where

- $a(u)f(u) + b(u)g(u) = m$;
- f is decreasing and g is increasing.

If $\mu \preceq_{DC} \nu$, it is straightforward to establish that $\nu_{[0,u]} = S^{\nu}(\mu_{[0,u]}) = S^{\nu}(u\delta_m)$, so that $(\nu_u)_{u\in[0,1]}$ is of the same type as $(\mu_u)_{u\in[0,1]}$, that is

$$\nu_u = a'(u)\delta_{f'(u)} + b'(u)\delta_{g'(u)}.$$

Moreover $f' \le f$ and $g \le g'$ and π_u is concentrated on the four oriented pairs $(f(u), f'(u))$, $(g(u), g'(u))$, $(f(u), g'(u))$ and $(f'(u), g(u))$. More explicitely.

$$\pi_u = \frac{1}{g'(u) - f'(u)} \Big([(g'(u) - f(u))\delta_{f(u),f'(u)} + (f(u) - f'(u))\delta_{f(u),g'(u)}] \\ + [(g'(u) - g(u))\delta_{g(u),f'(u)} + (g(u) - f'(u))\delta_{g(u),g'(u)}] \Big).$$

Finally note that if $(\mu_t)_{t\in T}$ is a family of probability measures indexed by a partial order T so that $s \le t$ implies $\mu_s \preceq_{DC} \mu_t$, then there exist a martingale $(X_t)_{t\in T}$ with $\mathrm{Law}(X_t) = \mu_t$ for every $t \in T$, [20, Theorem 4].

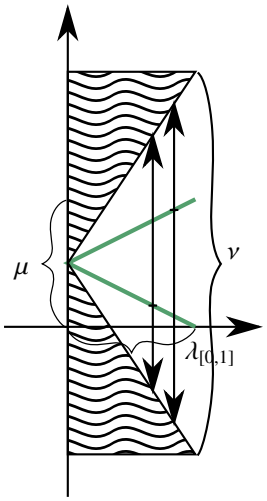

FIGURE 4. Middle-curtain coupling of uniform measures

3.1.4. *Comparison with the stochastic order, the quantile and independent couplings.* The stochastic order and the convex order share several common features and a parallel presentation can be given for shadow couplings of measures in convex order and the classical couplings. As already explained in [6] the left-curtain coupling can be considerd as the natural counterpart of the quantile coupling. Let us see that in the construction of the left-curtain coupling, replacing the shadows by *stochastic shadows* we obtain the quantile coupling. We assume $\mu \preceq_{\mathrm{sto}} \nu$. Let the stochastic shadow of a (sub)measure $\mu' \preceq_+ \mu$ in ν be the smallest measure η in $\preceq_{\mathrm{sto}}$ such that $\mu' \preceq_{\mathrm{sto}} \eta$ and $\eta \preceq_+ \nu$. The lift $\hat{\mu}$ is the quantile coupling of λ and μ, so that $\mu_{[0,u]} = (G_\mu)_\# \lambda|_{[0,u]}$. For $\mu \preceq_{\mathrm{sto}} \nu$, the stochastic shadow of this measure is simply $(G_\nu)_\# \lambda|_{[0,u]}$ independently of μ. Therefore, for the curves $(\mu_{[0,u]})_u$ and $(\nu_{[0,u]})$ defined in this way we recognize the quantile coupling.

We relax now the first condition $\mu_{[0,u]} \preceq_{\mathrm{sto}} \eta$ for defining the stochastic shadow (it is always satisfied for measures of mass u with $\eta \preceq_+ \nu$), so that we can establish a coupling not only in the case $\mu \preceq_{\mathrm{sto}} \nu$ but in general, keeping $\nu_{[0,u]} = (G_\nu)_\# \lambda|_{[0,u]}$. We obtain again the quantile coupling. The left-curtain and the quantile coupling are also analogous on the level optimality properties, see [6, Sections 1.2, 1.3].

While the left-curtain coupling can be viewed as the quantile coupling of the convex order world, we will explain next in which sense the sunset coupling corresponds to the independent (aka product) coupling. As before we do not assume $\mu \preceq_{\mathrm{sto}} \nu$ and we still

consider stochastic shadows given through $\nu_{[0,u]} = (G_\nu)_\# \lambda|_{[0,u]}$. However we take the same lift $\hat\mu = \lambda \times \mu$ as for the sunset coupling, so that $\mu_u = \mu$. It is then easy to identify the derivative in the target space as $\nu_u = \delta_{G_\nu(u)}$. Therefore the kernel final coupling writes

$$\pi_{[0,1]} = \pi = \int_0^1 \mu(\mathrm{Id} \times P_{G_\nu(u)})\, \mathrm{d}u.$$

This is nothing but $\mu \times \nu$. Apart from the Kellerer-Choquet representation given in Paragraph 3.1.2, we have encountered another sense in which the sunset coupling is particularly canonical: It is the product coupling of the convex order world.

4. Proof of Theorem 1.1

Throughout this section we assume that μ, ν are in convex order and that $\hat\mu \in \Pi(\lambda, \mu)$. Given a measurable $c : [0,1] \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}_+$ we consider the optimization problem

$$P := P_c := \inf \left\{ \int c \, \mathrm{d}\hat\pi : \hat\pi \in \hat\Pi_M(\hat\mu, \nu) \right\}. \tag{8}$$

Proposition 4.1. *Assume that $\hat\pi$ is the lifted shadow coupling corresponding to a curve $(\pi_{[0,u]})_{u\in[0,1]}$ as in Theorem 1.1 (4). Then for all $p \in [0,1], q \in \mathbb{R}$, $\hat\pi$ is an optimizer of* (8) *for $c_{p,q}(u,x,y) = \mathbb{1}_{u\le p}|y-q|$.*

Proof. Let $\hat\pi \in \hat\Pi_M(\hat\mu, \nu)$ be a shadow coupling and (p,q) as in the statement. Then $\int c_{p,q}\, \mathrm{d}\hat\pi = \int |y-q|\, \mathrm{d}\nu_{[0,p]}(y)$ where we recall $\nu_{[0,p]} = S^\nu(\mu_{[0,p]})$. More generally if $\hat\gamma$ is an element of $\hat\Pi_M(\hat\mu, \nu)$ we have $\int c_{p,q} \mathrm{d}\hat\gamma = \int |y-q|\, \mathrm{d}\beta^p(y)$ where $\mu_{[0,p]} \preceq_C \beta^p$ and $\beta^p \preceq_+ \nu$ (in fact $\beta^p := (\mathrm{proj}_y)_\# \gamma|_{[0,p]\times\mathbb{R}\times\mathbb{R}}$). Therefore $\nu_{[0,q]} \preceq_C \beta^p$ and as $y \mapsto |y-q|$ is convex, we have proved that $\hat\pi$ is a minimizer. □

Actually if $\hat\gamma$ is a minimizer, for any $p \in [0,1]$ the measures β^p and $\nu_{[0,p]}$ have the same potential function. Thus they are equal and the curve $p \mapsto (\mathrm{proj}_y)_\# \gamma|_{[0,p]\times\mathbb{R}\times\mathbb{R}})$ is completely determined. But we have proved the uniqueness of such couplings in Theorem 1.1. Hence $\hat\gamma = \hat\pi$.

Recall from Theorem 1.1 that a set $\hat\Gamma \subseteq [0,1] \times \mathbb{R} \times \mathbb{R}$ is called *monotone* if for all $u, v, x, x', y^-, y^+, y'$ such that $u < v, (u, x, y^-), (u, x, y^+), (v, x', y') \in \Gamma$ it holds $y' \notin]y^-, y^+[$.

Proposition 4.2. *Assume that $\hat\pi \in \Pi_M(\mu, \nu)$ satisfies one of the following assumptions:*

(1) For all $u \in [0,1], q \in \mathbb{R}$, $\hat\pi$ is an optimizer of (8) *for $c_{p,q}(u,x,y) = \mathbb{1}_{u\le p}|y-q|$.*

(2) $\hat\pi$ is an optimizer of (8) *for $c(u,x,y) = (1-u)\sqrt{1+y^2}$.*

Then there is a monotone set $\hat\Gamma$ such that $\hat\pi(\hat\Gamma) = 1$.

Proof. We will establish the assertion under the first assumption, the argument based on the second assumption is very similar.

Using the notation and the monotoneity principle from [2], we pick for each $(u,q) \in ([0,1] \times \mathbb{R}) \cap \mathbb{Q}^2$ a monotoneity set $\Gamma_{(u,q)}$ for the cost function and set

$$\Gamma := \bigcap_{(u,q)\in[0,1]\times\mathbb{R}\cap\mathbb{Q}^2} \Gamma_{(u,q)}.$$

Assume for contradiction that there exist $s, t, x, x', y^-, y^+, y'$ such that

$$s < t, a := (s, x, y^-), b := (s, x, y^+), c := (t, x', y,) \in \Gamma \Rightarrow y' \in]y^-, y^+[.$$

Pick λ such that $y' = (1-\lambda)y^- + \lambda y^+$, $u \in]s,t[$, and q very close to y' (in comparison to y^-, y^+). Set

$$a' := (t, x', y^-), b := (t, x', y^+), c := (s, x, y,).$$

Then

$$\alpha := (1-\lambda)\delta_a + \lambda\delta_b + \delta_c, \alpha' := (1-\lambda)\delta_{a'} + \lambda\delta_{b'} + \delta_{c'}$$

are competitors with $\operatorname{supp}\alpha \subseteq \hat{\Gamma}$ and $\int c_{u,q}\,\mathrm{d}\alpha > \int c_{u,q}\,\mathrm{d}\alpha'$, contradiction. □

In the next result we establish that any $\hat{\pi}$ which is monotone admits a barrier representation as in Theorem 1.1 (2).

Proposition 4.3. *Let $\hat{\pi} \in \Pi_M(\hat{\mu}, \nu)$ be a transport plan concentrated on a monotone set $\hat{\Gamma}$. Define barriers*

$$R_o := \{(s,y) \in [0,1]\times\mathbb{R} : \exists t > s, (t,y) \in \hat{\Gamma}\} \tag{9}$$

$$R_c := \{(s,y) \in [0,1]\times\mathbb{R} : \exists t \geq s, (t,y) \in \hat{\Gamma}\}. \tag{10}$$

Consider a process $(Z_t)_{t\geq 0} = (Z_t^1, Z_t^2)_{t\geq 0} = (Z_0^1, Z_t^2)_{t\geq 0}$ on some probability space which takes values in $[0,1]\times\mathbb{R}$ and is specified through

(1) $Z_0 \sim \hat{\mu}$,

(2) $Z_t = Z_0 + (0, B_t)$, where $(B_t)_t$ is (one dimensional) Brownian motion.

and write τ_o, τ_c for the first time Z hits R_o resp. R_c. Then $\tau_o = \tau_c$ a.s. and

$$(Z_0, Z_{\tau_o}) \sim (Z_0, Z_{\tau_c}) \sim \hat{\pi}. \tag{11}$$

The martingales $t \mapsto Z_{t\wedge\tau}$ and $t \mapsto B_{t\wedge\tau}$ are uniformly integrable.

There exist Borel maps $T_{up}, T_{down} : [0,1]\times\mathbb{R} \to \mathbb{R}, T_{down}(x) \leq x \leq T_{up}(x)$ such that

$$\hat{\pi}\{(u,x,T_i(x)) : i \in \{up, down\}, (u,x) \in [0,1]\times\mathbb{R}\} = 1.$$

Proof. Fix a disintegration $(\pi_{u,x})$ of $\hat{\pi}$ wrt $\hat{\mu}$ and write $\Gamma_{u,x}$ for the section of $\hat{\Gamma}$ in (u,x). Then $\hat{\mu}(\Gamma_0) = 1$, where

$$\Gamma_0 = \left\{(u,x) : \pi_{u,x}(\Gamma_{u,x}) = 1, \int |y|\,\mathrm{d}\pi_{u,x} < \infty, \int y\,\mathrm{d}\pi_{u,x} = x\right\}.$$

Define for each $(u,x) \in [0,1]\times\mathbb{R}$, $\tau_{u,x}$ to be the Azema-Yor solution (say) of the Skorokhod embedding problem such that $B_{\tau_{u,x}} \sim \pi_{u,x}$. Then define a stopping time τ such that conditionally on $Z_0 = (u,x)$ we have $\tau = \tau_{u,x}$. It follows that $(Z_0, Z_\tau^2) \sim \hat{\pi}$ and that for all elements ω of a full measure set Ω_0 we have $(Z_0^1(\omega), Z_0^2(\omega), Z_\tau^2(\omega)) \in \hat{\Gamma}$. Next we claim that there exists a full measure subset Ω_1 of Ω_0 such that for all $\omega \in \Omega_1$ and every $t < \tau(\omega)$, the following assertion holds true:

Continuation Assertion on (ω, t). There are $\omega_i \in \Omega_0, i = 1,2$ satisfying

(1) $t < \tau(\omega_i)$, $(Z_s(\omega))_{s\leq t} = (Z_s(\omega_i))_{s\leq t}$ for $i = 1,2$,

(2) $Z_\tau^2(\omega_1) < Z_t^2(\omega) < Z_\tau^2(\omega_2)$.

Assume for contradiction that the set

$$\{(\omega,t) : t < \tau(\omega) \text{ and Continuation Assertion fails}\} =: D$$

is not evanescent, i.e. that $\mathrm{proj}_\Omega(D)$ does not have $\mathbb{P}$-measure 0. Set

$$D^- := \{(\omega,t) \in [\![0,\tau[\![: \omega_1 \in \Omega_0, t < \tau(\omega_1), (Z_s(\omega_1))_{s\leq t} = (Z_s(\omega))_{s\leq t} \Rightarrow Z_\tau(\omega_1) \leq Z_t(\omega)\}$$

$$D^+ := \{(\omega,t) \in [\![0,\tau[\![: \omega_2 \in \Omega_0, r < \tau(\omega_2), (Z_s(\omega_2))_{s\leq t} = (Z_s(\omega_2))_{s\leq t} \Rightarrow Z_\tau(\omega_2) \geq Z_t(\omega)\}$$

such that $D = D^- \cup D^+$. If D is not evanescent, then by the optional section theorem there exists a stopping time σ such that $\mathbb{P}(\sigma < \infty) > 0$ and

$$\{(\omega, \sigma(\omega)) : \sigma(\omega) < \infty\} \subseteq D^- \text{ or } \{(\omega, \sigma(\omega)) : \sigma(\omega) < \infty\} \subseteq D^+.$$

Combined with the strong Markov property this leads to a contradiction with the optional stopping theorem.

We claim that on Ω_1

$$\tau_c \leq \tau \leq \tau_o. \tag{12}$$

Note that the first inequality is satisfied by definition of τ_c. To establish the second inequality we assume for contradiction that there exists $\omega \in \Omega_1$ such that $\tau_o(\omega) < \tau(\omega)$.

Then $t^* := \min\{t \geq 0 : Z_t(\omega) \in R_o\} < \tau(\omega)$. Set $y' := Z^2_{t^*}(\omega)$ and $(u, x) = (Z^1_0(\omega), Z^2_0(\omega))$. By definition of R_o, there exist $v > u$ and x' such that $(v, x', y') \in \hat{\Gamma}$. Pick $\omega_i, i = 1, 2$ according to the Continuation Assertion. Setting $y_i = Z^2_\tau(\omega_i), i = 1, 2$, we have $(u, x, y_i) \in \hat{\Gamma}$, contradiction.

By Lemma 4.4 $\tau_c = \tau_o$ almost surely hence (11) holds.

To see that $(Z_{t\wedge\tau})$ (resp. $(B_{t\wedge\tau})$) is uniformly integrable we recall a result of Monroe [25] which asserts that a solution τ of the Skorokhod problem is minimal (i.e. there is no strictly smaller solution) if and only if Brownian motion up to time τ is uniformly integrable. In the present context it is straight forward to verify that τ provides a minimal embedding of ν wrt Z^2 (we refer to [4, Proposition 4.1] for complete details), hence $(Z^2_{t\wedge\tau})$ is uniformly integrable.

The rest is immediate. □

In the proof we used the following lemma from [8] (we include the proof for the convenience of the reader).

Lemma 4.4. *Let $\hat{\mu}$ be a probability measure on $\mathbb{R}^2$ such that the projection onto the horizontal axis* $\mathrm{proj}_x \hat{\mu}$ *is continuous (in the sense of not having atoms) and let $\phi : \mathbb{R} \to \mathbb{R}$ be a Borel function. Set*

$$R_o := \{(x, y) : x > \phi(y)\}, \quad R_c := \{(x, y) : x \geq \phi(y)\}.$$

Start a vertically moving Brownian motion in μ and define

$$\tau_o := \inf\{t : (x, y + B_t) \in R_o\}, \quad \tau_c := \inf\{t : (x, y + B_t) \in R_c\}.$$

Then $\tau_c = \tau_o$ almost surely.

Proof. Obviously $\tau_c \leq \tau_o$.

We say that y is a local minimum of ϕ if $\phi(y') \geq \phi(y)$ for all y' in a neighborhood of y. Set

$$I := \{\phi(y) : y \text{ is a local minimum of } \phi\}.$$

It is then not difficult to prove (and certainly well known) that I is at most countable: assume by contradiction that there exist an uncountably family $A \subseteq \mathbb{R}$ and corresponding neighborhoods $(a - \varepsilon_{sun}, a + \varepsilon_a), a \in A$ such that $\phi(x) \geq \phi(a)$ for $x \in (a - \varepsilon_a, a + \varepsilon_a)$ and $a \neq a'$ implies $f(a) \neq f(a')$. Passing to an uncountable subset of A, we can assume that there is some $\eta > 0$ such that $\varepsilon_a > \eta$ for all $a \in A$. For $a \neq a'$ we cannot have $|a - a'| < \eta$ for then $a \in (a' - \varepsilon_{a'}, a' + \varepsilon_{a'})$ as well as $a' \in (a - \varepsilon_a, a + \varepsilon_a)$ which would imply that $f(a) = f(a')$. Hence $|a - a'| \geq \eta$ which implies that A is countable, giving a contradiction.

On the complement of $I \times \mathbb{R}$ we have almost surely

$$\tau_o = 0 \quad \Longleftrightarrow \quad \tau_c = 0 \tag{13}$$

as a consequence of the strong Markov property. □

We have thus obtained an interpretation of monotone transport plans in terms of a barrier-type solution to the Skorokhod problem. This interpretation is useful for us since it allows us to use a short argument of Loynes [24] (which in turn builds on Root [26]) to show that there is only one monotone transference plan.

Lemma 4.5 (cf. Loynes [24])**.** *Let $\hat{\pi}_1$, $\hat{\pi}_2$ be monotone transport plans in $\hat{\Pi}_M(\hat{\mu}, \nu)$, with corresponding maps $T^{\mathbb{P}} = (T^i_{up}, T^i_{down})$ and denote by $R^{\hat{\pi}_i}, i = 1, 2$ the corresponding 'closed' barriers as in Proposition 4.3. Then $\tau_{R^{\hat{\pi}_1}} = \tau_{R^{\hat{\pi}_2}}$, a.s.*

Proof. For a set $A \subseteq \mathbb{R}$, we abbreviate $R_i(A) := R^{\hat{\pi}_i} \cap (\mathbb{R} \times A)$ and $\tau_i = \tau_{R^{\hat{\pi}_i}}$ for $i = 1, 2$. Denote

$$K := \{y : m_1(y) > m_2(y)\} \text{ where } \qquad m_i(y) := \sup\{m : (m, y) \in R^{\hat{\pi}_i}\},\ i = 1, 2. \tag{14}$$

Fix a trajectory $(Z_t)_t = (Z_t(\omega))_t$ such that $Z^2_{\tau_2} \in K$. Then $(Z_t)_t$ hits $R_2(K)$ before it enters $R_2(K^C)$. But then $(Z)_t$ also hits $R_1(K)$ before it enters $R_1(K^C)$. Hence

$$B_{\tau_2} \in K \quad \Longrightarrow \quad B_{\tau_1} \in K.$$

As both stopping times embed the same measure, this implication is an equivalence almost surely, and we may set $\Omega_K := \{B_{\tau_1} \in K\} = \{B_{\tau_2} \in K\}$. On Ω_K we have $\tau_1 \leq \tau_2$ while $\tau_1 \geq \tau_2$ on Ω^C_K. Then, for all Borel subset $A \subseteq \mathbb{R}$:

$$\mathbb{P}[B_{\tau_1 \wedge \tau_2} \in A] = \mathbb{P}[B_{\tau_1 \wedge \tau_2} \in A, \Omega_K] + \mathbb{P}[B_{\tau_1 \wedge \tau_2} \in A, \Omega^c_K] \tag{15}$$

$$= \mathbb{P}[B_{\tau_1} \in A, \Omega_K] + \mathbb{P}[B_{\tau_2} \in A, \Omega^c_K] \tag{16}$$

$$= \mathbb{P}[B_{\tau_1} \in A \cap K] + \mathbb{P}[B_{\tau_2} \in A \cap K^c] \tag{17}$$

$$= \mathbb{P}[B_{\tau_2} \in A \cap K] + \mathbb{P}[B_{\tau_2} \in A \cap K^c] \tag{18}$$

$$= \mathbb{P}[B_{\tau_2} \in A] \tag{19}$$

since $B_{\tau_i} \sim \nu$. Hence $\tau_1 \wedge \tau_2$ embeds ν. Similarly, we see that $\tau_1 \vee \tau_2$ also embeds ν. Since τ_1 and τ_2 are both minimal embeddings, we deduce that $\tau_1 \wedge \tau_2 = \tau_1$ as well as $\tau_1 \wedge \tau_2 = \tau_2$. □

Taking the results of this section we can now establish our main theorem.

Proof of Theorem 1.1. We have already seen in Theorem 3.1 that there exists $\hat{\pi} \in \hat{\Pi}_M(\hat{\mu}, \nu)$ satisfying Theorem 1.1 (4). By virtue of Propositions 4.1, 4.2 we have that $\hat{\pi}$ is monotone as required in 1.1 (3) and by Proposition 4.3 $\hat{\pi}$ admits a barrier type representation as in 1.1 (2). Moreover, by Lemma 4.4 we find that there exists a unique such $\hat{\pi}$.

Finally, by the standard compactness-continuity argument there exists $\hat{\pi}$ which solves the optimization problem in Theorem 1.1 (1), by Proposition 4.2 it is monotone and hence uniquely determined as before. □

Note that in the proof of Theorem 1.1, we did not use the uniqueness part in the statement of Theorem 3.1; rather we have obtained a second derivation of this uniqueness property based on Lemma 4.5.

We close this section with a remark on the implication of the above results for the curtain coupling.

Remark 4.6. We consider the curtain coupling π_{lc} corresponding to the case where the lift $\hat{\mu}$ is given by the monotone rearrangement between Lebesgue measure and μ. Assume for simplicity that μ has no atoms such that the lift $\hat{\mu}$ is concentrated on the graph of a 1-1 function elements of $\Pi_M(\mu, \nu)$ correspond in a 1-1 manner to elements of $\hat{\Pi}_M(\hat{\mu}, \nu)$. It then follows from the respective optimality property of $\hat{\pi}$ that π_{lc} minimizes

$$\gamma \mapsto \int \phi(x)\psi(y)\, \mathrm{d}\gamma(x, y) \tag{20}$$

on the set $\Pi_M(\mu,\nu)$, where $\varphi \geq 0$ is strictly decreasing and $\psi \geq 0$ is strictly convex and the minimum over $\Pi_M(\mu,\nu)$ is finite. Moreover, there exist a Borel set $S \subseteq \mathbb{R}$ and two measurable functions $T_1, T_2 : S \to \mathbb{R}$ such that

(1) π_{lc} is concentrated on the graphs of T_1 and T_2.
(2) For all $x \in \mathbb{R}$, $T_1(x) \leq x \leq T_2(x)$.
(3) For all $x < x' \in \mathbb{R}$, $T_2(x) < T_2(x')$ and $T_1(x') \notin]T_1(x), T_2(x)[$.

This recovers [6, Corollary 1.6].

5. The sunset coupling as a non-optimizer and shadow couplings as optimizers to general transport problems.

An important message of [6, 14] is that the left-curtain couplings are characterized as the optimizers to martingale optimal transport problems for a large class of cost functions. This goes together with the fact that the support of the left-curtain coupling is typically a very 'small' set – if μ is continuous, it is contained in the graphs of two functions.

In contrast, the sunset coupling typically has a 'large' support. Hence we do not expect it to solve a martingale transport problem except in trivial instances. This is underlined by the following simple example.

Example 5.1. Let μ, ν be measures in convex order such that

(1) $\mathrm{conv}(\mathrm{supp}(\mu)) \cap \mathrm{supp}(\nu) = \emptyset$
(2) μ, ν consist of finitely many atoms.

Assume c is such that the sunset coupling is optimal for the martingale transport problem. Then *all* elements of $\Pi_M(\mu,\nu)$ are optimal for the martingale transport problem.

Proof. We first note that $\mathrm{supp}(\pi_{\mathrm{sun}}) = \mathrm{supp}(\mu) \times \mathrm{supp}(\nu)$ under our assumptions on μ, ν.

In the present atomic case, the martingale transport problem can be formulated as a linear programming problem and it admits a natural dual problem for which strong duality holds. It follows from this that every martingale transport plan $\pi \in \Pi_M(\mu,\nu)$ satisfying $\mathrm{supp}(\pi) \subseteq \mathrm{supp}(\pi_{\mathrm{sun}})$ is optimal. □

For simplicity, we have stated Example 5.1 for discrete marginals but (with some work) it is not difficult to see that the same phenomenon carries over to more general cases.

However we find it interesting to note that shadow couplings posses optimality properties in a different sense: Recently, Gozlan et al. [13] introduced a framework for *general* transport problems. As in the classical case one optimizes over the set of transport plans $\pi \in \mathbb{P}(\mu,\nu)$, where μ, ν are probabilities on Polish spaces X, Y. In contrast to the classical case, more general cost functionals are considered. Writing $\mathcal{P}(Y)$ for the set of all probability measures on Y, a general cost is a function $C : X \times \mathcal{P}(Y) \to [0,\infty]$ and its associated transport costs are

$$T_C(\nu|\mu) := \inf_{\pi \in \Pi(\mu,\nu)} \int C(x, \bar{\pi}_x)\, \mathrm{d}\mu(x), \tag{21}$$

where we use $(\bar{\pi}_x)_x$ to denote disintegration wrt μ.

The shadow couplings appear as optimizers to such general transport problems. E.g. we will see below that the sunset coupling is the unique optimizer for the general transport cost function

$$C(x, \bar{\pi}) := \inf_{\alpha \in \hat{\Pi}_M(\lambda \times \delta_x, \bar{\pi})} \int (1-u)\sqrt{1+y^2}\, \mathrm{d}\alpha(u, x', y). \tag{22}$$

Here the function $(u, y) \mapsto (1-u)\sqrt{1+y^2}$ could be replaced by any function of the form $(u, y) \mapsto \phi(u)\psi(y)$, where ϕ is strictly increasing and ψ strictly convex and sufficiently integrable wrt the given marginals. The cost function defined in (22) exhibits a relatively intuitive behavior: If $\bar{\pi}$ does not have center x the costs equal $+\infty$. If $\bar{\pi}$ is centered around x, the more $\bar{\pi}$ is spread out, the higher are the costs.

Proposition 5.2. *Let μ and ν be in convex order and fix a disintegration $(\bar{\mu}_x)_x$ of $\hat{\mu}$ wrt μ and set*

$$C^{\hat{\mu}}(x, \bar{\pi}) := \inf_{\alpha \in \hat{\Pi}_M(\bar{\mu}_x \times \delta_x, \bar{\pi})} \int (1-u)\sqrt{1+y^2}\, \mathrm{d}\alpha(u, x', y). \tag{23}$$

Then the shadow coupling associated to $\hat{\mu}$ is the unique optimizer of the general transport problem associated to $C^{\hat{\mu}}$.

We note that the solution to the optimization problems (22) / (23) is straightforward to characterize in the nontrivial case where x is the barycenter of $\bar{\pi}$: The optimizer α is the unique element of $\hat{\Pi}_M(\bar{\mu}_x \times \delta_x, \bar{\pi})$ which is concentrated on the graphs of two functions $T_{\mathrm{up}} : [0,1] \to [x, \infty)$, $T_{\mathrm{down}} : [0,1] \to (-\infty, x]$, where T_{up} is increasing and T_{down} is decreasing.

Proof. Given $\hat{\pi} \in \hat{\Pi}_M(\hat{\mu}, \nu)$, write π for the corresponding martingale transport $\pi \in \Pi_M(\mu, \nu)$ and $(\bar{\pi}_x)_x$ for its disintegration wrt μ. We note that $\hat{\pi}$ can be μ-a.s. uniquely represented in the form

$$\hat{\pi}(A \times B \times C) = \int \mathrm{d}\mu(x) \int \mathrm{d}\alpha_x(u, x', z)\, \mathbb{1}_{A \times B \times C}(u, x, y), \tag{24}$$

where (α_x) is a (measurable) family with $\alpha_x \in \hat{\Pi}_M(\bar{\mu}_x \times \delta_x, \bar{\pi}_x)$. We then find

$$\begin{aligned} \inf_{\hat{\pi} \in \hat{\Pi}_M(\hat{\mu}, \nu)} \int (1-u)\sqrt{1+y^2}\, \mathrm{d}\hat{\pi} &= \inf_{\pi \in \Pi_M(\mu, \nu)} \inf_{\alpha \in \hat{\Pi}_M(\bar{\mu}_x \times \delta_x, \bar{\pi}_x)} \iint (1-u)\sqrt{1+y^2}\, \mathrm{d}\alpha(u, x', y)\, \mathrm{d}\mu(x) \\ &= \inf_{\pi \in \Pi(\mu, \nu)} \int C^{\hat{\mu}}(x, \bar{\pi}_x)\, \mathrm{d}\mu(x). \end{aligned} \tag{25}$$

□

References

[1] M. Beiglböck, A. Cox, and M. Huesmann. Optimal transport and skorokhod embedding. *Invent. Math., to appear*, 2016.

[2] M. Beiglböck and C. Griessler. An optimality principle with applications in optimal transport. *ArXiv e-prints*, Apr. 2014.

[3] M. Beiglböck, P. Henry-Labordère, and F. Penkner. Model-independent bounds for option prices: A mass transport approach. *Finance and Stochastics*, 17(3):477–501, 2013.

[4] M. Beiglböck, P. Henry-Labordère, and N. Touzi. Monotone Martingale Transport Plans and Skorokhod Embedding. *preprint*, 2015.

[5] M. Beiglböck, M. Huesmann, and F. Stebegg. Root to Kellerer. *Séminaire de Probabilités, to appear*, 2016.

[6] M. Beiglböck and N. Juillet. On a problem of optimal transport under marginal martingale constraints. *Ann. Probab.*, 44(1):42–106, 2016.

[7] M. Beiglböck, M. Nutz, and N. Touzi. Complete Duality for Martingale Optimal Transport on the Line. *Ann. Probab., to appear*, 2015.

[8] M. Beiglböeck, A. M. G. Cox, and M. Huesmann. The geometry of multi-marginal Skorokhod Embedding. *preprint*, 2016.

[9] B. Bouchard and M. Nutz. Arbitrage and duality in nondominated discrete-time models. *The Annals of Applied Probability*, 25(2):823–859, 2015.

[10] L. Campi, I. Laachir, and C. Martini. Change of numeraire in the two-marginals martingale transport problem. *ArXiv e-prints*, June 2014.

[11] Y. Dolinsky and H. M. Soner. Martingale optimal transport and robust hedging in continuous time. *Probab. Theory Relat. Fields*, 160(1-2):391–427, 2014.
[12] A. Galichon, P. Henry-Labordère, and N. Touzi. A Stochastic Control Approach to No-Arbitrage Bounds Given Marginals, with an Application to Lookback Options. *Ann. Appl. Probab.*, 24(1):312–336, 2014.
[13] N. Gozlan, C. Roberto, P.-M. Samson, and P. Tetali. Kantorovich duality for general transport costs and applications. *ArXiv e-prints*, Dec. 2014.
[14] P. Henry-Labord'ere and N. Touzi. An Explicit Martingale Version of Brenier's Theorem. *Finance Stoch., to appear*, 2016.
[15] F. Hirsch, C. Profeta, B. Roynette, and M. Yor. *Peacocks and associated martingales, with explicit constructions*, volume 3 of *Bocconi & Springer Series*. Springer, Milan, 2011.
[16] F. Hirsch and B. Roynette. A new proof of Kellerer's theorem. *ESAIM Probab. Stat.*, 16:48–60, 2012.
[17] D. Hobson and M. Klimmek. Robust price bounds for the forward starting straddle. *ArXiv*, Apr. 2013.
[18] D. Hobson and A. Neuberger. Robust bounds for forward start options. *Mathematical Finance*, 22(1):31–56, 2012.
[19] N. Juillet. Stability of the shadow projection and the left-curtain coupling. *to appear in Ann. Inst. Henri Poincaré Probab. Stat.*, 2014.
[20] N. Juillet. Peacocks parametrised by a partially ordered set. In *Séminaire de Probabilités, to appear*. 2016.
[21] H. Kellerer. Markov-Komposition und eine Anwendung auf Martingale. *Math. Ann.*, 198:99–122, 1972.
[22] H. G. Kellerer. Integraldarstellung von Dilationen. In *Transactions of the Sixth Prague Conference on Information Theory, Statistical Decision Functions, Random Processes (Tech. Univ., Prague, 1971; dedicated to the memory of Antonín Špaček)*, pages 341–374. Academia, Prague, 1973.
[23] G. Lowther. Fitting Martingales To Given Marginals. *ArXiv e-prints*, Aug. 2008.
[24] R. M. Loynes. Stopping times on Brownian motion: Some properties of Root's construction. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 16:211–218, 1970.
[25] I. Monroe. On embedding right continuous martingales in Brownian motion. *Ann. Math. Statist.*, 43:1293–1311, 1972.
[26] R. Rockafellar. Characterization of the subdifferentials of convex functions. *Pacific J. Math.*, 17:497–510, 1966.
[27] F. Stebegg. Model-Independent Pricing of Asian Options via Optimal Martingale Transport. *ArXiv e-prints*, Dec. 2014.
[28] V. Strassen. The existence of probability measures with given marginals. *Ann. Math. Statist.*, 36:423–439, 1965.
[29] C. Villani. *Optimal transport*, volume 338 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, 2009.